Measuring IPv6 Deployment

Geoff Huston
APNIC
December 2009

IPv4 address exhaustion

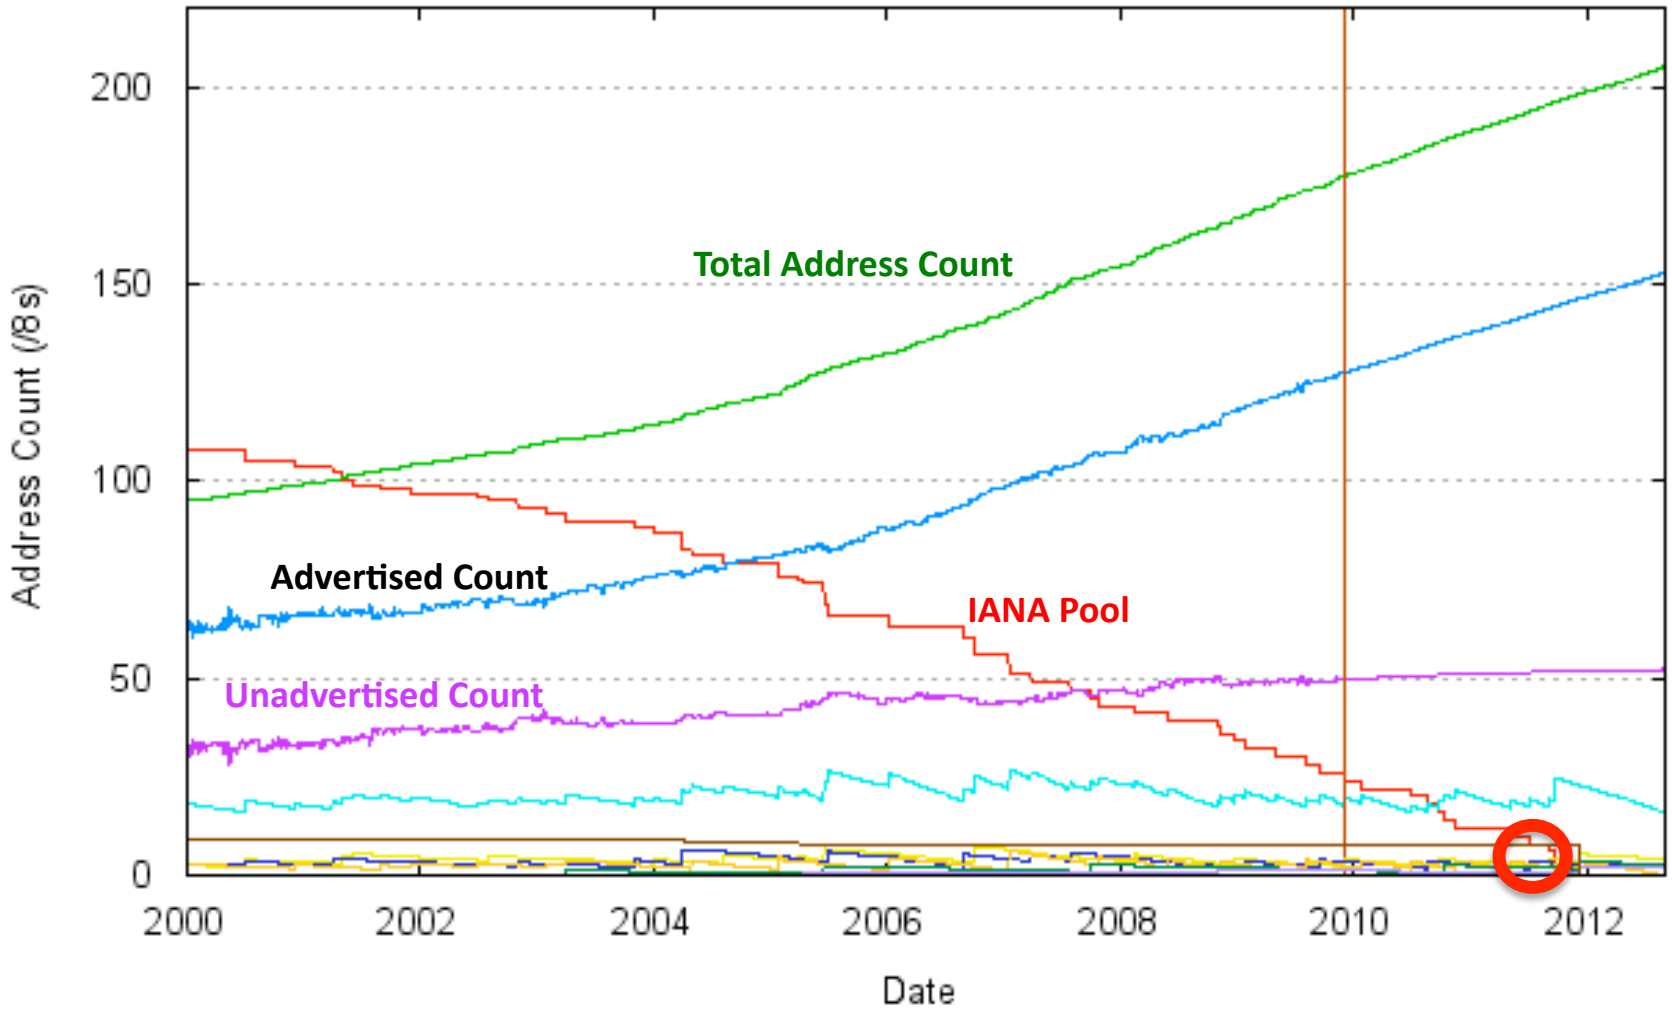

IPv4 address exhaustion

The model of address consumption predicts that IANA will allocate its last ‘general use /8 address block on the 5th September 2011

The first RIR to exhaust its “general use” address pool will be APNIC, on the 30th August 2012

http://ipv4.potaroo.net

The dual stack transition plan

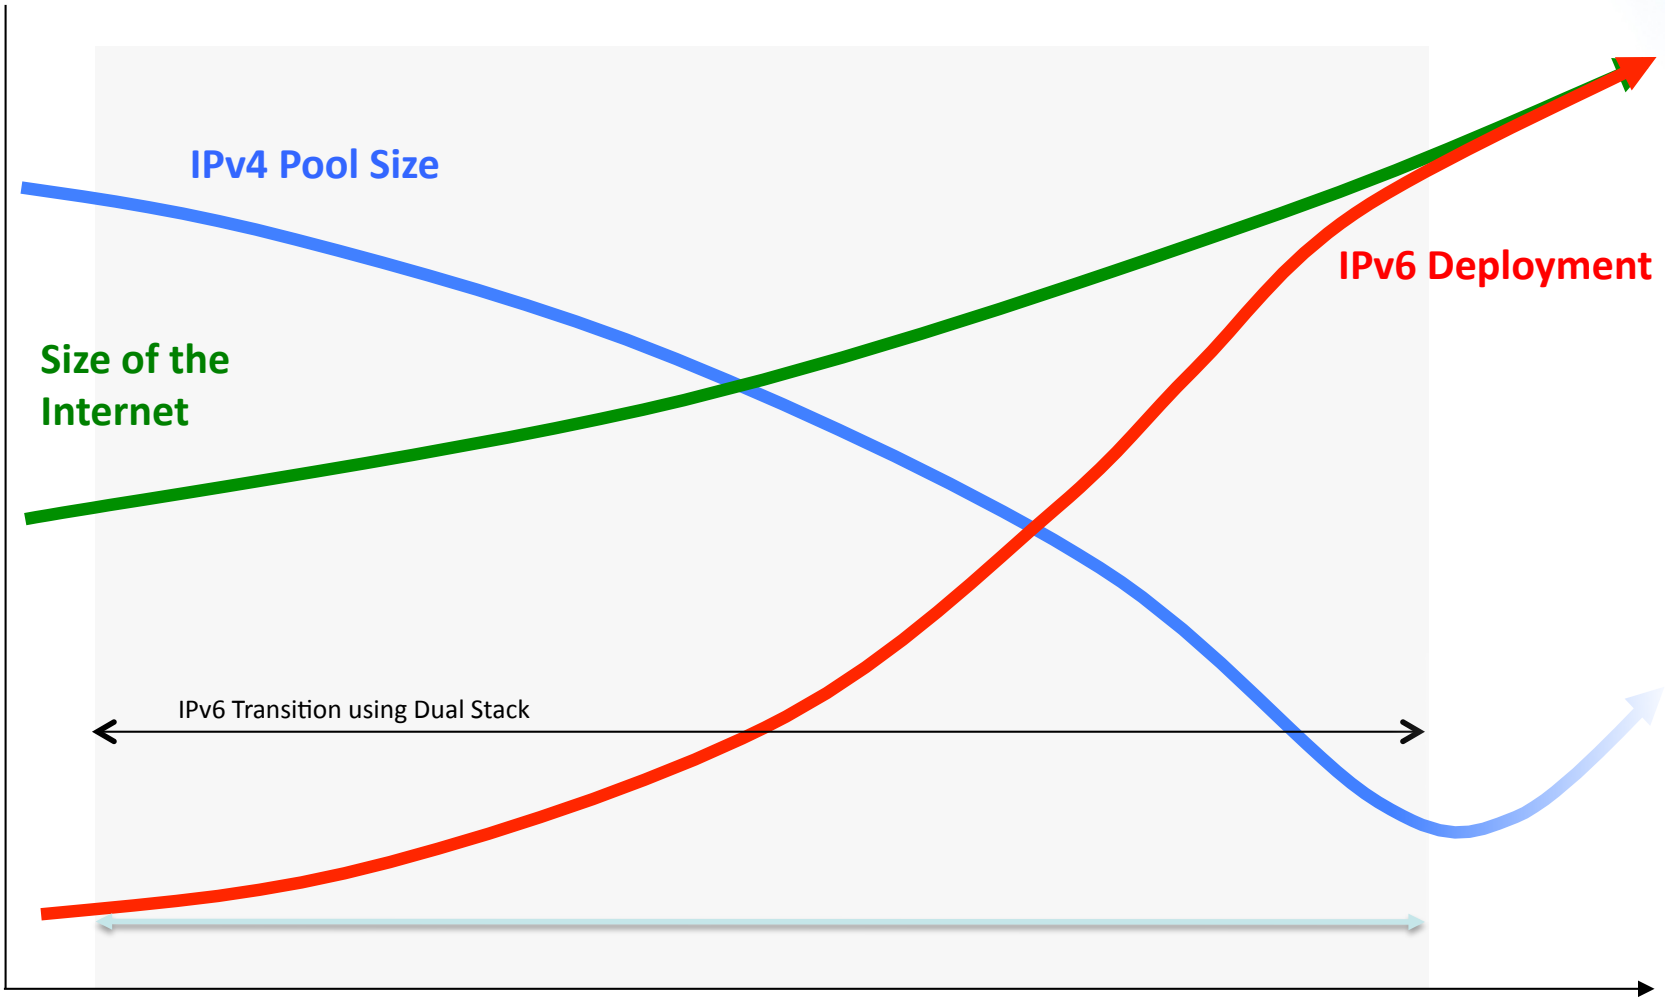

Where are we with IPv6 transition?

What long-term network metric data sets can indicate the pace of relative deployment of IPv6?

Measurement issues

- Measurement approaches:
  - Whole of network metrics vs sample measurements
  - Component metrics vs system metrics
  - Static metrics vs performance metrics
  - Snapshot metrics vs time series measurements

- The nature of the data set vs the nature of the measurement
  - Is the data set distribution heavy-tail or Gaussian?
  - Is the measurement reflective of end-user behaviour or infrastructure?
  - Is the time series fractal or converging?

What's the question?

Candidate questions:

- How much of the public Internet supports v6?
- How much of the public Internet runs v6?
- How quickly is the Internet becoming end-to-end v6 capable?
- How long will the dual stack transition take?

I. Whole of network measurements

- There is only one obvious “whole of network” vantage point for the Internet: Inter-Domain Routing, or BGP
  - Even this view is limited to transit paths
- BGP Question:
  - How ‘large’ is the v6 BGP routing space compared to the v4 BGP routing space?

IPv6 routing table entries

IPv4 routing table entries

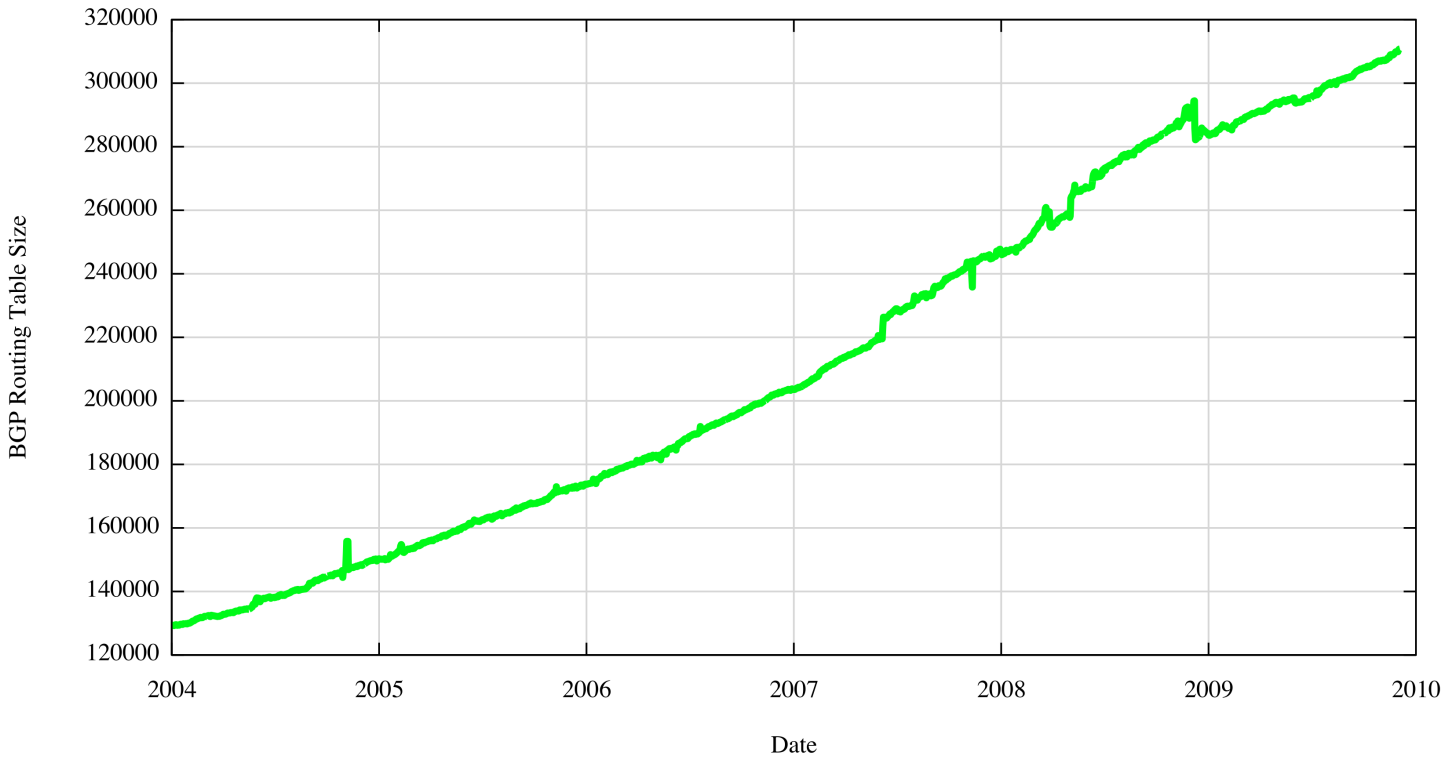

Combined view: IPv6 and IPv4

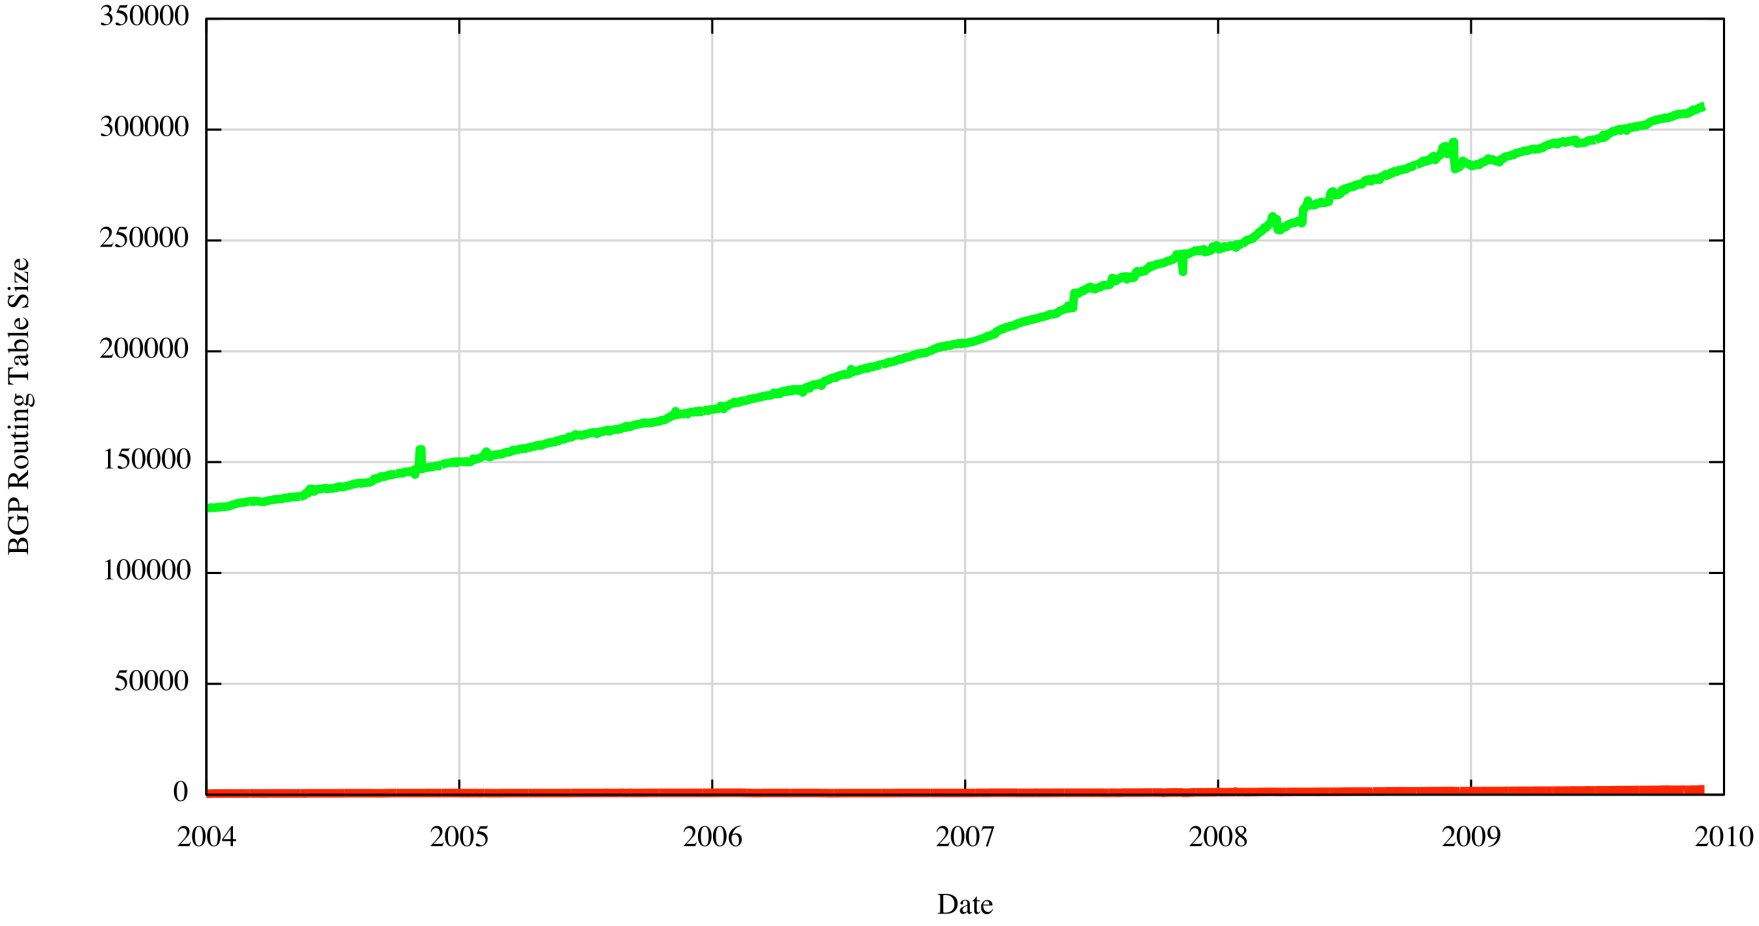

Ratio of IPv6 to IPv4

What's this saying?

- IPv6 is currently **0.8%** of IPv4 in terms of routing table entries
- Assuming future exponential growth of this ratio, v6 will be at 80% of the v4 Internet in 2026
  - This projection is extremely unreliable. The change of conditions following IPv4 free pool exhaustion will significantly alter the dynamics of IPv6 uptake

Is this a good indicator of IPv6 deployment?

Probably not!

The two data sets are not directly comparable. Why?

- Historical fragmentation in IPv4
- Traffic engineering in IPv4
- Sparsely used announcements in IPv6
- Use of tunneling prefixes in IPv6 without routing entries

Redefining the question

How much of the network is capable of supporting IPv6?

- A possible metric for this question in BGP can be found in the number of ASes that announce and/or transit IPv6 addresses

IPv6 AS count

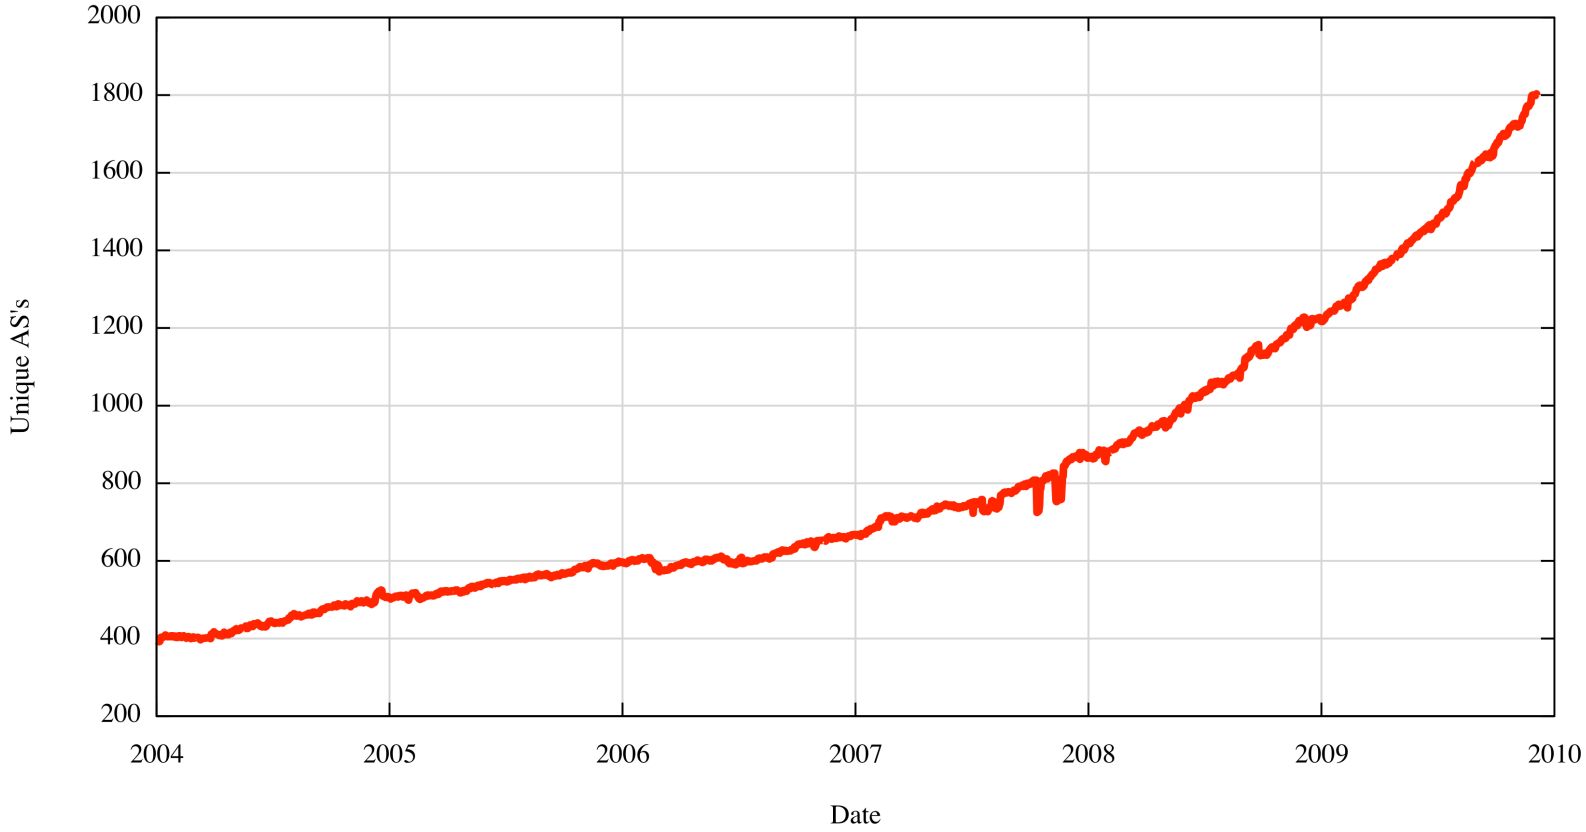

IPv4 AS count

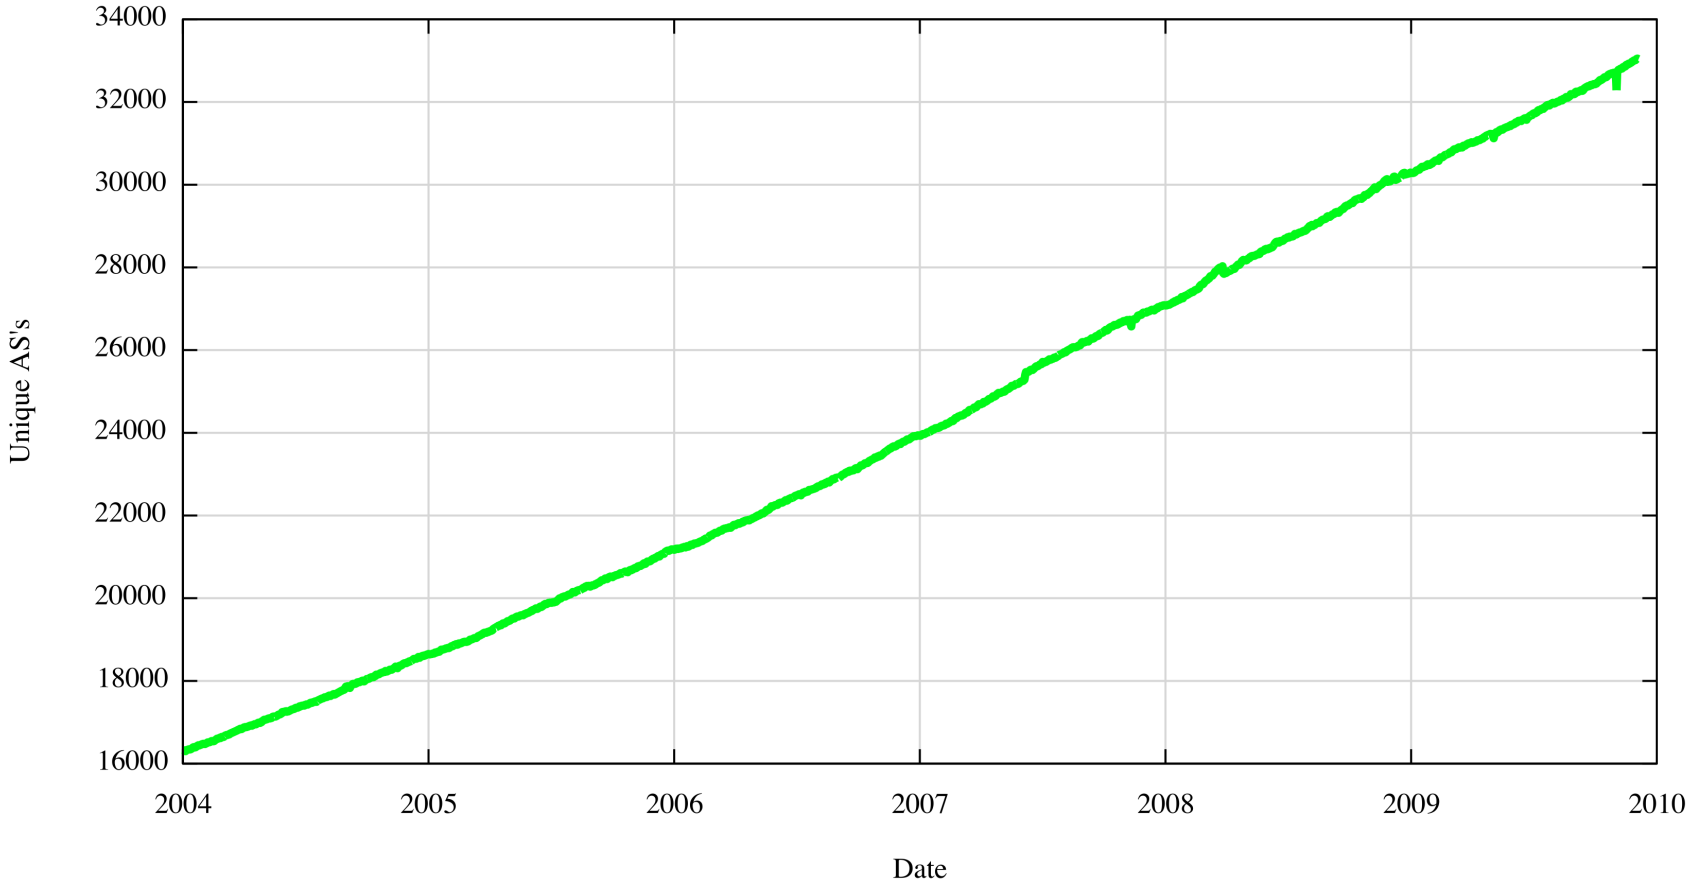

Ratio of IPv6 to IPv4 ASes

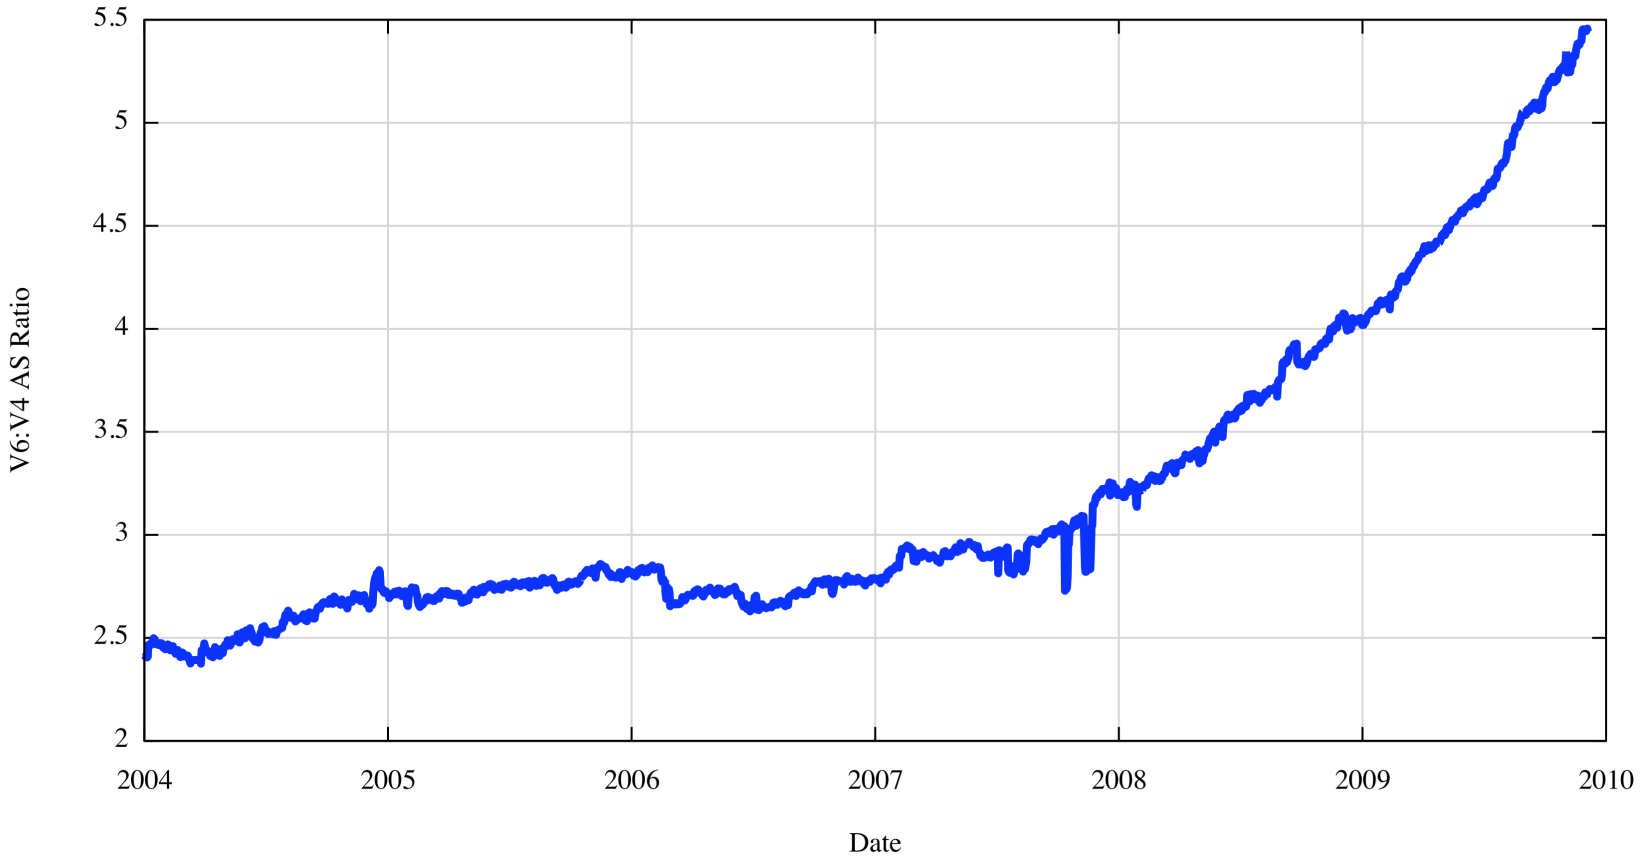

What's this saying?

- IPv6 is currently **5.5%** of IPv4 in terms of ASes that announce or transit IPv6 routes
- Assuming future exponential growth of this ratio, v6 will be at 80% of the v4 Internet in 2018

Transit ASes vs stub ASes

- This 5.5% is not uniformly distributed across all ASes
- Of the 33,039 ASes in IPv4, 28,596 ASes are ‘stub’ networks (non transit) and 4,443 are transit
  - Of these 4,443 v4 **transit** ASes, 910 also announce v6 prefixes (**20%**) - double the value of March 2009
  - of the 28,596 v4 **stub** ASes, 887 also announce v6 prefixes (**3%**)
- Transit SPs in IPv4 are very active in IPv6, and an 80% deployment level by 2013 is a reasonable projection from these numbers

Interpretation of BGP data

- Routing table metrics are not a clear indicator of actual levels of uptake given the disparity of routing structures between IPv4 and IPv4
- Transit AS uptake appears to be a possible indicator of IPv6 uptake in the ISP sector
  - But, this metric does not expose the level of deployment of customer services for IPv6. Advertising an IPv6 prefix in BGP is not the same as providing IPv6 services to attached end sites.
- Perhaps this approach can be used for other indirect measurements of IPv6 deployment
  - Is the IPv6 inter-AS topology similar to the IPv4 BGP network? Or is the IPv6 network still being constructed as an overlay using a different set of connectivity enablers and constraints?

II. Infrastructure data

- This is part of a set of metrics that measure individual components of network infrastructure for IPv6 support
  - Potential Measurements include
    - Counts of domain names with AAAA records
    - Counts of domain name servers with AAAA records
    - Samples of queries for AAAA records from servers
  - APNIC data set
    - Queries for reverse mapping of DNS names to IPv4 / IPv6 addresses

DNS data

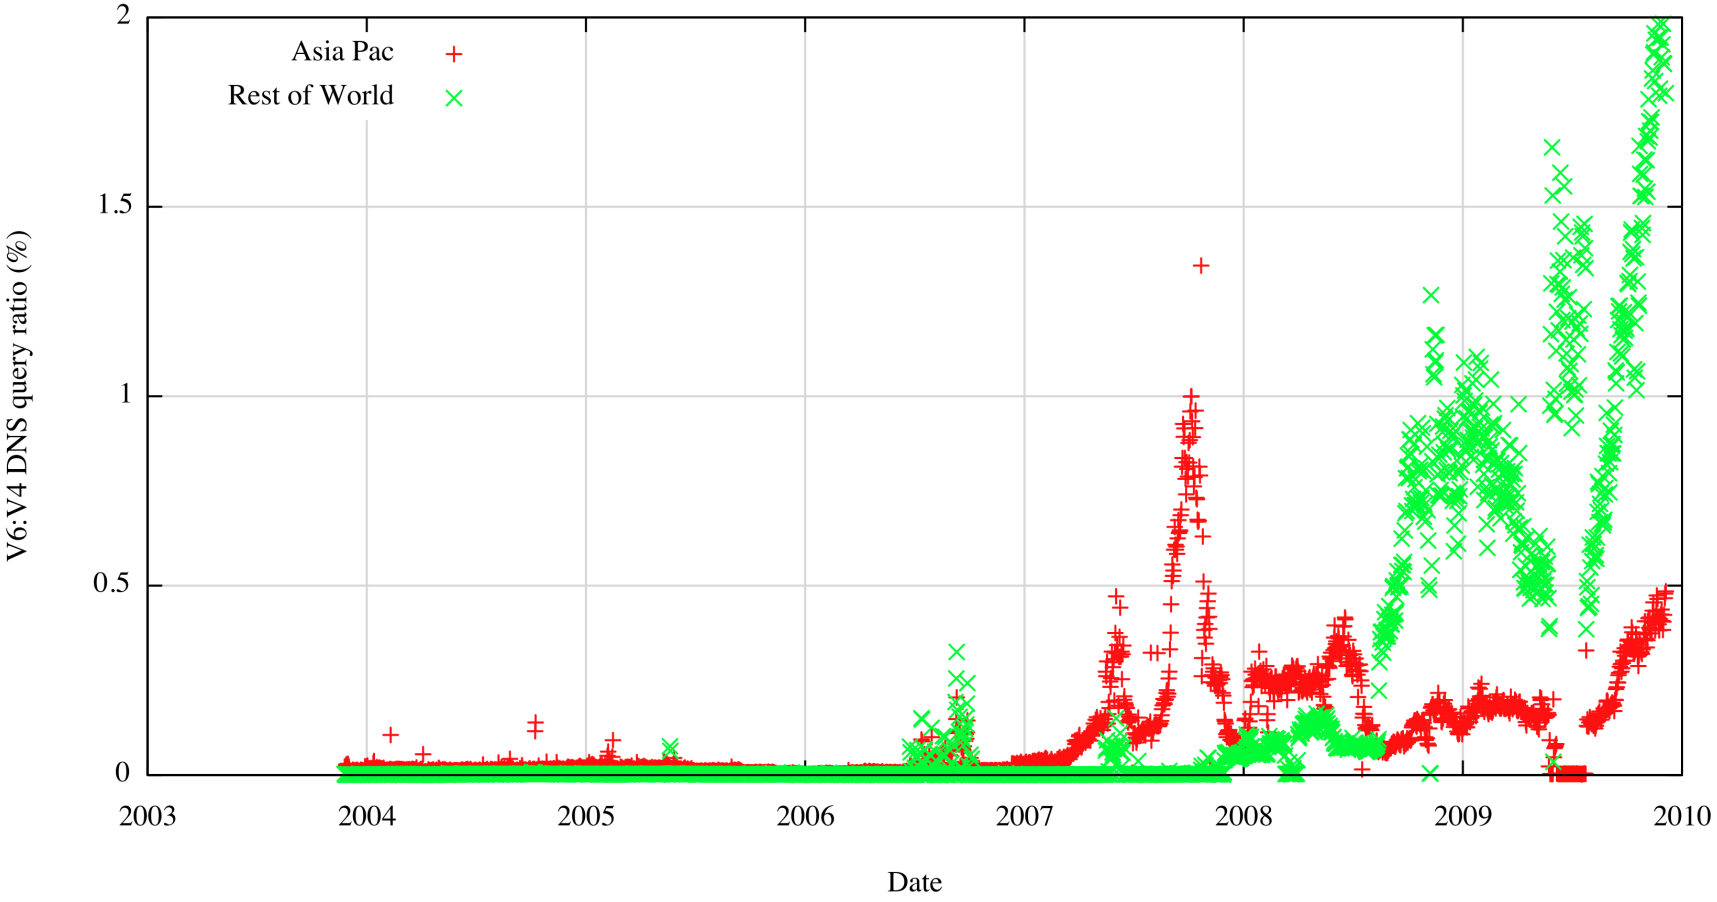

DNS data

- This is difficult data to interpret because of the interactions between DNS forwarders and DNS caches and query behaviours for reverse DNS resolution
- The data appears to be broadly reflective of IPv6 deployment levels, but the relationship is very indirect and there are many variables in this data
- The high variance between resolvers for different address sets is anomalous

III. End-to-end service data

- Examine IPv6 / IPv4 use from the perspective of a service delivery platform (web server)
- IPv6 is selected by clients only when all the various IPv6 infrastructure components support IPv6, otherwise the client will fall back to IPv4 use
- Service metrics for IPv6 are reflective of end-to-end IPv6 capability
- Simple sampling approach that any dual stack web server can use

Web server stats

- Take a couple of dual-homed web servers:
  - http://www.apnic.net
  - http://www.ripe.net
- Count the number of distinct IPv4 and IPv6 addresses per day
  - Not the number of web ‘hits’, just the ratio of the populations of distinct source addresses that access these sites, to reduce the relative impact of robots and crawlers on the data and normalize the data against different profiles of use
- Look at the v6 / v4 access ratio

APNIC web server stats

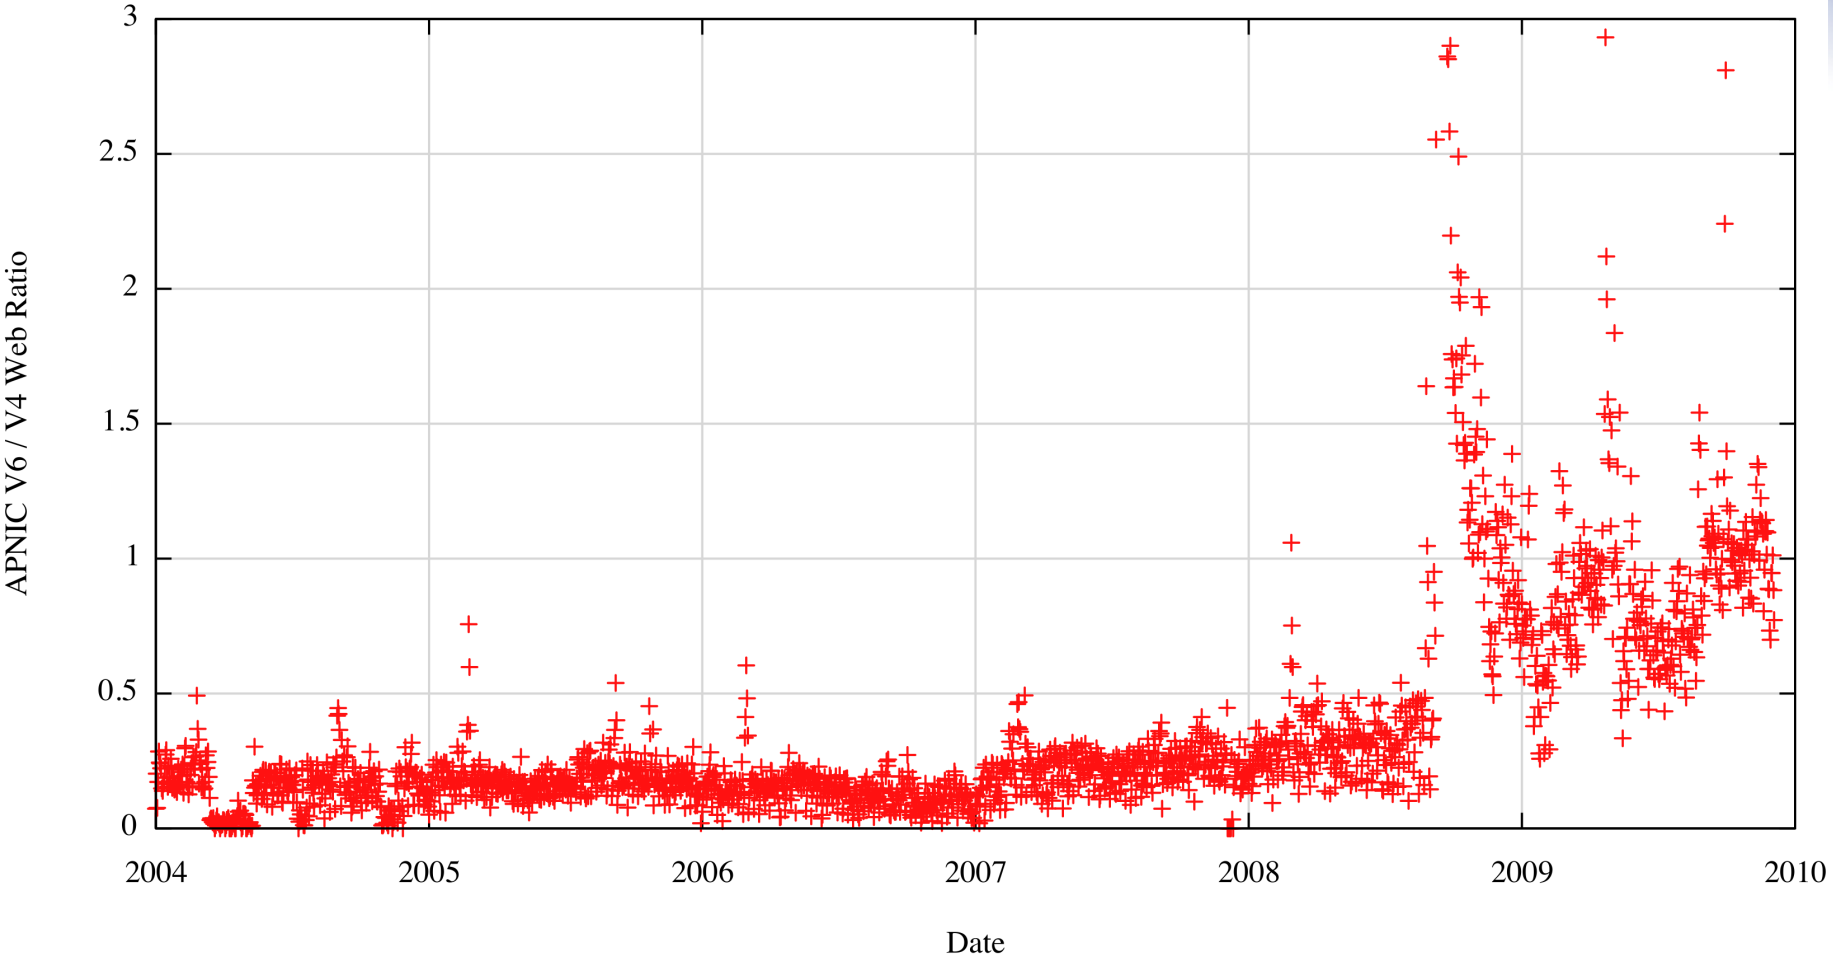

RIPE NCC web server stats

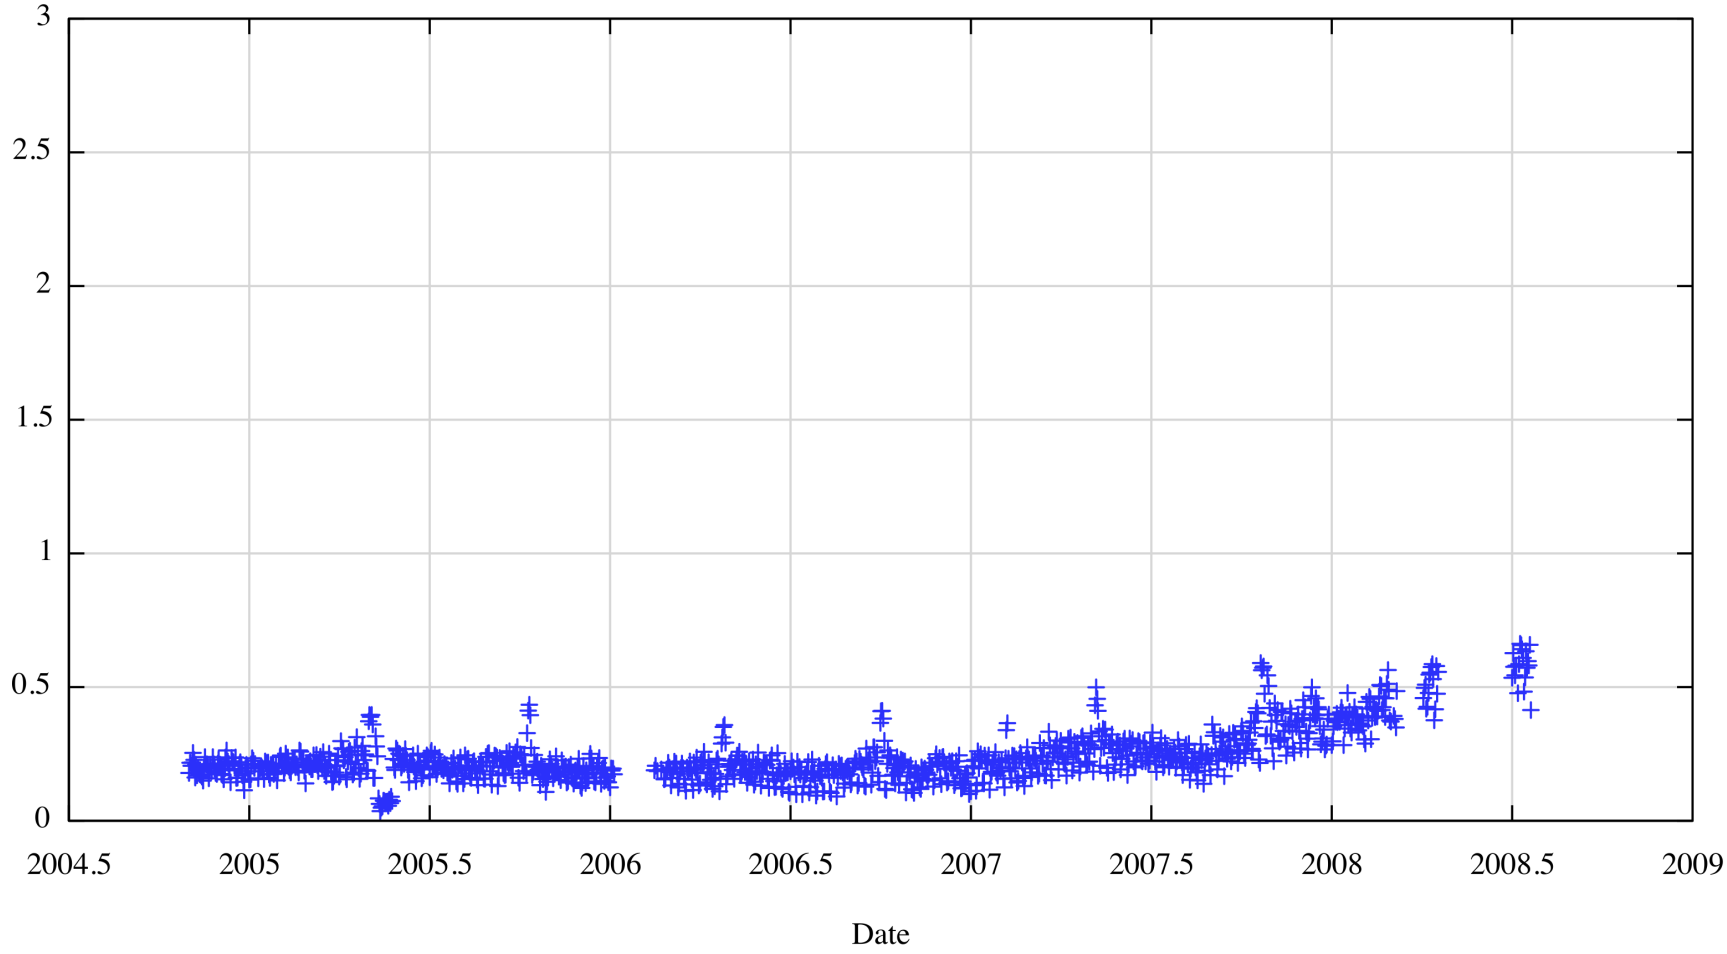

Combined stats

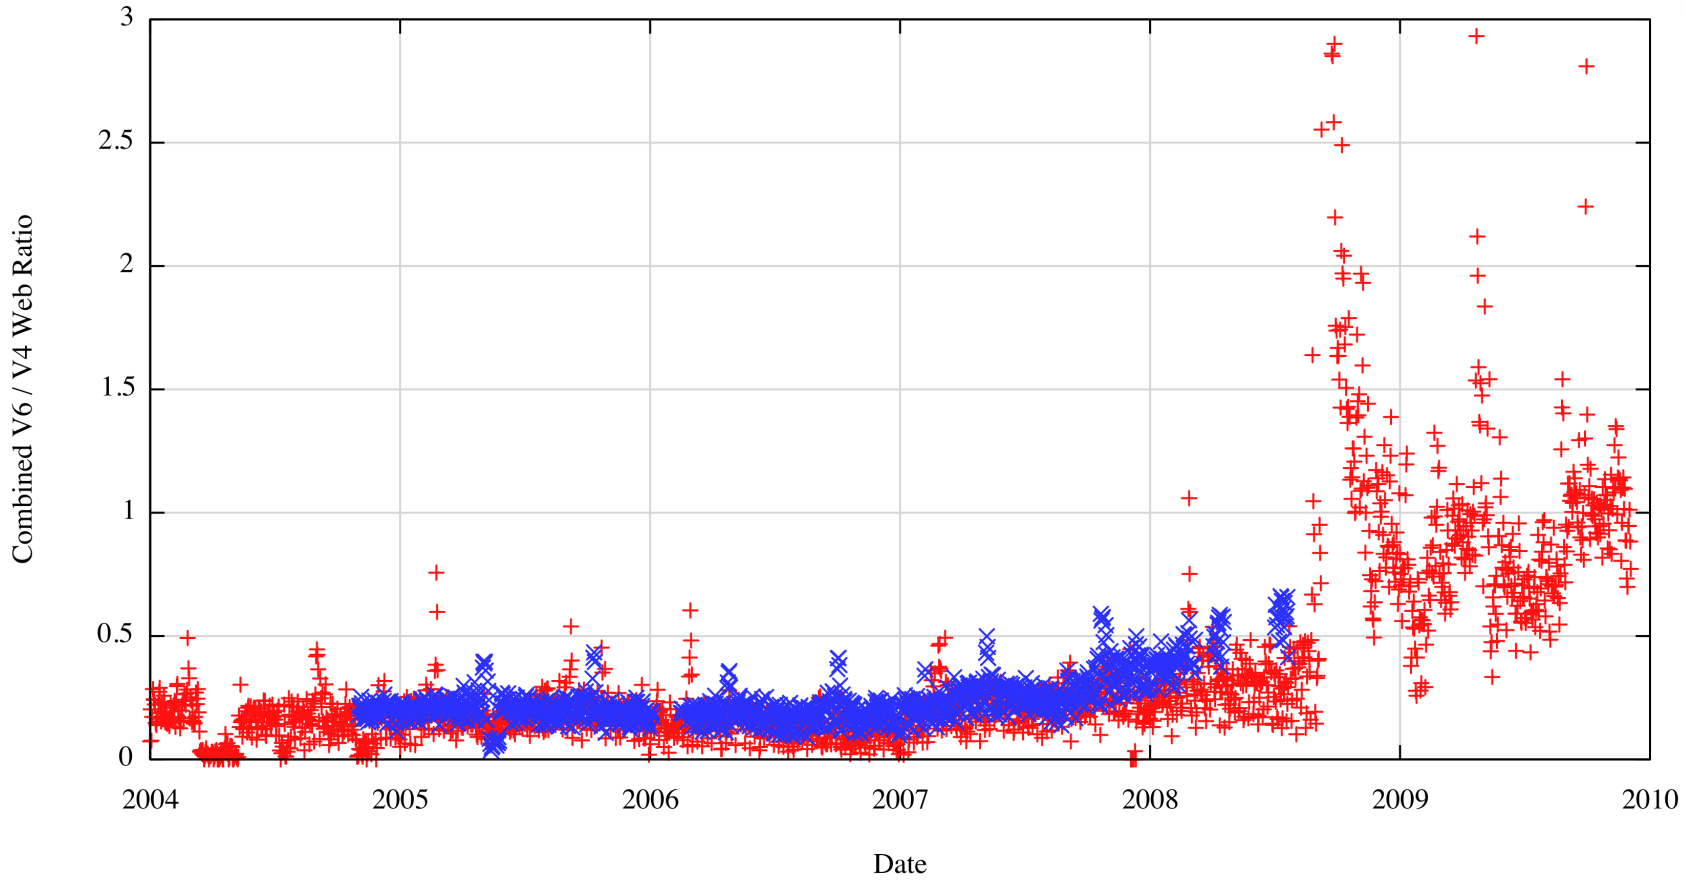

What's this saying?

Relative use of IPv6 has slowly increased over four years to reach around 1% today

How much of this is due to IPv6 over IPv4 tunneling?

IPv6 tunnelling

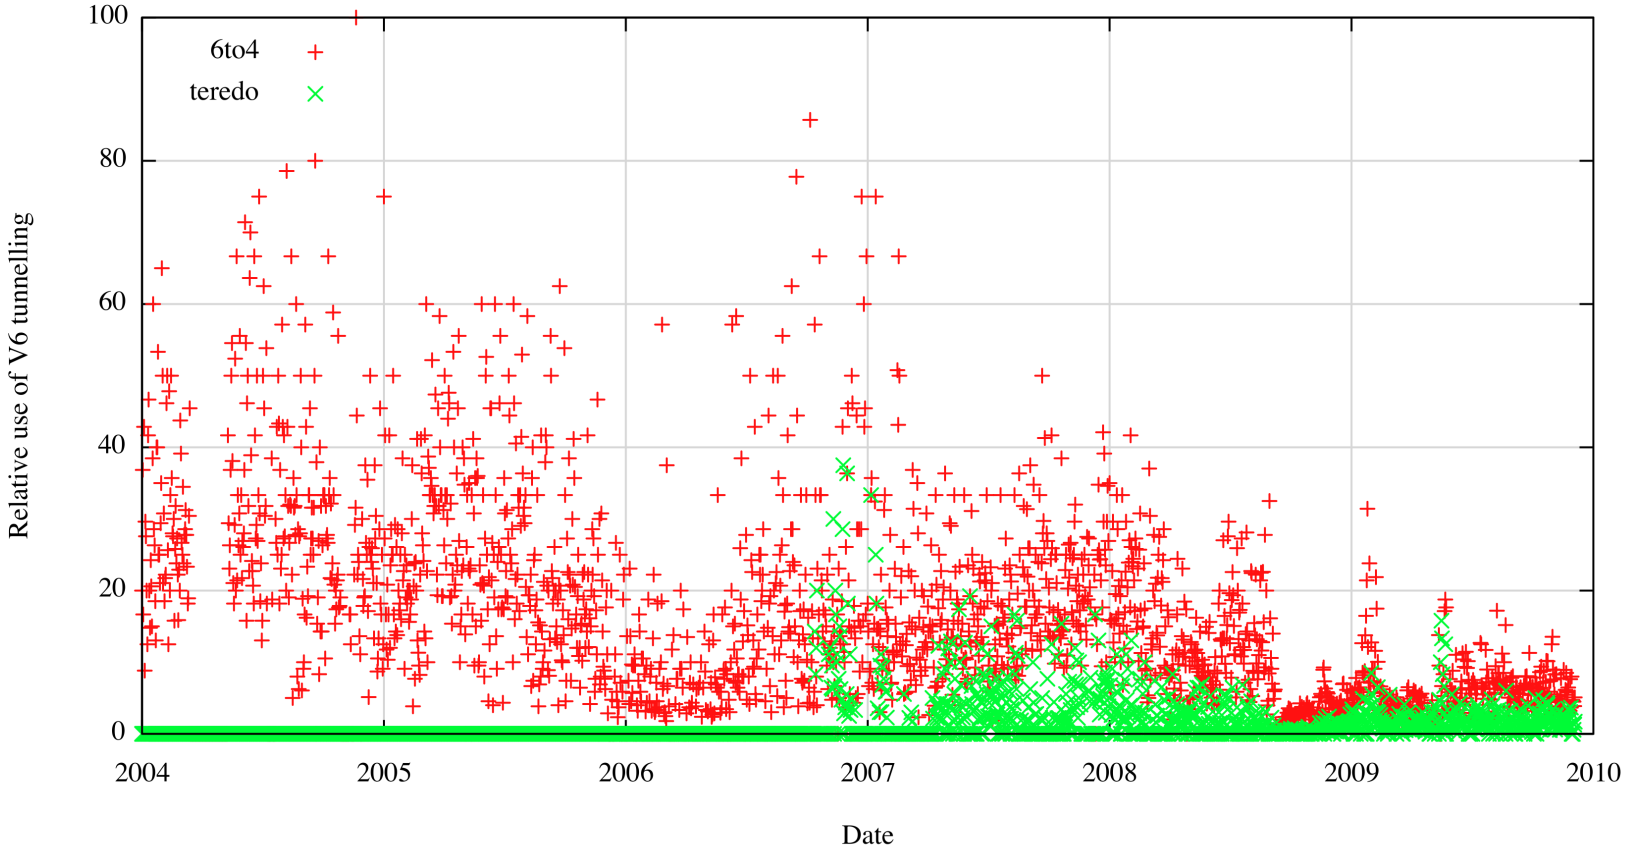

Where are we with IPv6?

The ‘size’ of the IPv6 deployment in terms of end-to-end host IPv6 capability is around 1% of the total number of Internet end hosts at present

Tunnelling represents around 10% of the IPv6 volume

Where are we with IPv6?

But, that 1% figure is probably too high

- Widespread NAT use in IPv4 undercounts IPv4 host counts
- These web sites are tech-weenie web sites. More general web sites and services probably have far fewer IPv6-capable clients
- Many mobile 3G-based hosts are still IPv4-only at present

Where are we with IPv6?

The current ‘size’ of the IPv6 deployment in terms of support for end-to-end client/ service transactions using IPv6 is somewhere in the range of 0.1% - 0.5% of the Internet’s end host population.

The current status of IPv6 deployment

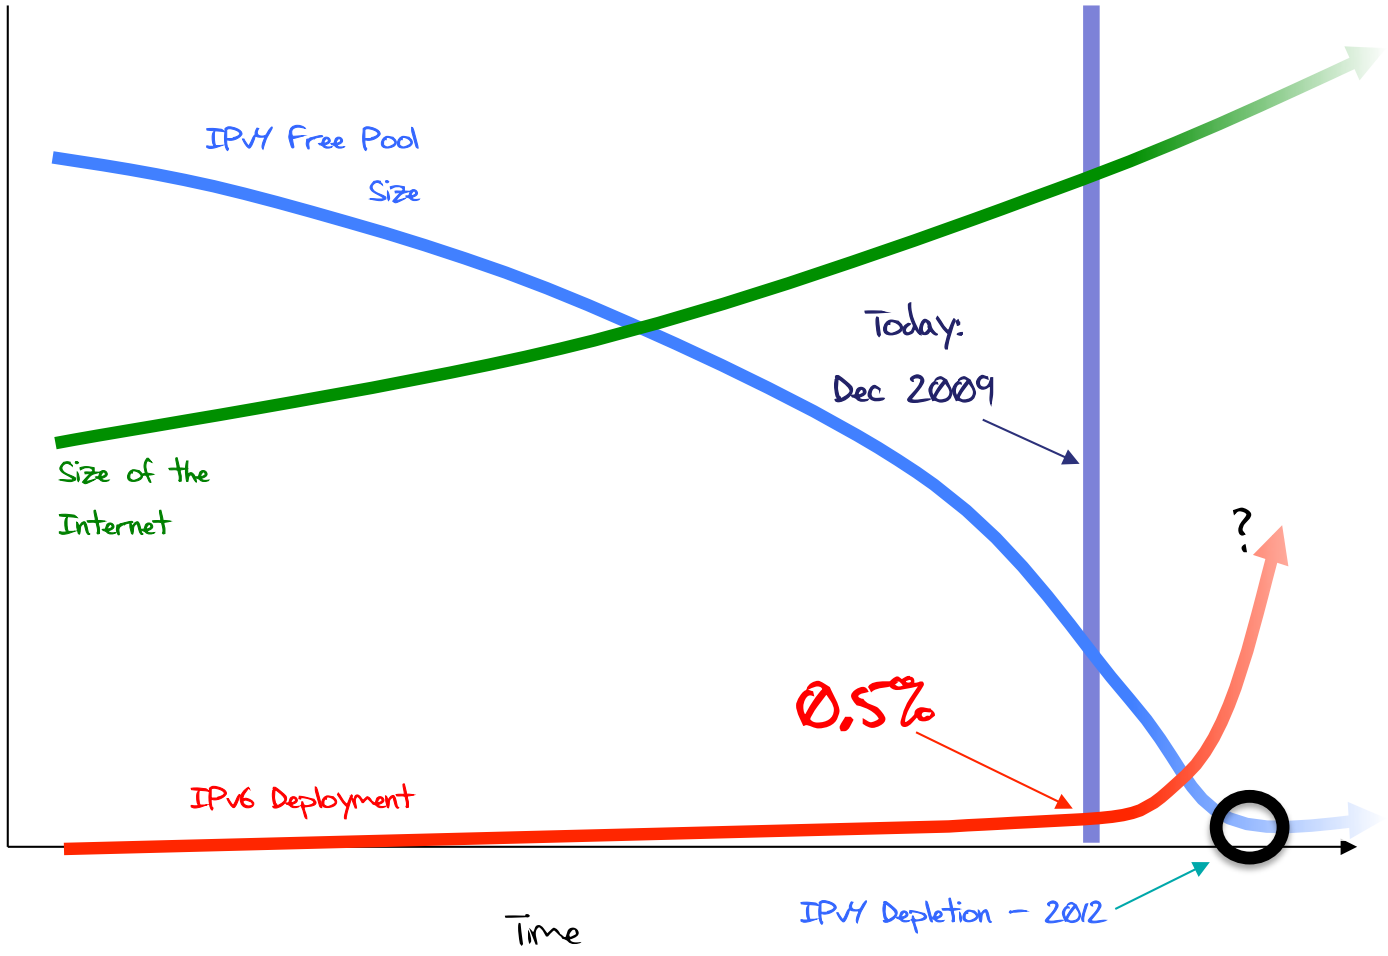

Thank you